

Veeva Vault RIM: AI for Submission Planning & Correspondence

By Adrien Laurent, CEO at IntuitionLabs • 2/13/2026 • 50 min read

veeva vault rim • submission planning • regulatory automation • generative ai in pharma • submission content plans
regulatory affairs



Executive Summary

In the highly regulated life-sciences industry, preparing and submitting dossiers to health authorities is a complex, document-intensive process. The work entails extensive **submission planning** (defining document content, schedules, and regulatory strategies) and managing **health-authority correspondence** (answering regulatory queries, meeting reports, and commitments). These tasks have traditionally been laborious and error-prone: teams use spreadsheets and ad hoc tools to track hundreds or thousands of submission documents and queries, often leading to delays and inconsistencies ^[1] www.veeva.com ^[2] www.freyrsolutions.com). However, modern **Regulatory Information Management** (RIM) systems like Veeva Vault RIM are consolidating regulatory data and documents onto a single cloud platform ^[3] www.veeva.com ^[4] www.veeva.com), enabling real-time visibility into submission status and query resolution. This report examines how **artificial intelligence (AI)** – particularly large language models and automation – can be integrated into the Veeva Vault RIM framework to **automate submission planning and health-authority correspondence**, dramatically accelerating workflows and improving quality.

Key findings from the analysis include:

- **Dramatic business impact of faster submissions:** Industry studies show that speeding up dossier filings yields enormous value. For a \$1 billion asset, filing just one month earlier can unlock roughly **\$60 million** in net present value ^[5] www.mckinsey.com). Leading biopharma companies are now filing **2–3x faster** than the industry average of 2020 ^[6] www.mckinsey.com), highlighting the critical importance of efficient regulatory processes. AI-powered automation is a major enabler of these gains. For example, McKinsey reports that an AI-assisted writing platform co-developed with Merck cut first-draft clinical study report (CSR) time from 180 to 80 hours ^[7] www.mckinsey.com). Similarly, practical implementations of intelligent automation have **reduced end-to-end authoring cycles by 40%** in trial documentation ^[8] www.mckinsey.com). Fast, accurate submission planning and query handling directly translate into earlier drug approvals and significant commercial benefit.
- **Veeva Vault RIM capabilities:** Veeva's Vault RIM suite (used by roughly 15 of the top 20 pharma companies ^[4] www.veeva.com)) provides content management tools for regulatory filings, including **Submission Content Plans**, **Health Authority Question (HAQ) tracking**, and automated publishing. For instance, Vault's Submission Content Plans module "auto-generates a table of contents" for a dossier and lets teams assign documents to the outline ^[9] www.veeva.com (regulatory.vaultbasics.help). The Vault platform also supports structured workflows: users can highlight text in a regulatory letter to create HAQ records, link those questions to a submission, and automatically generate response documents from templates (regulatory.veevavault.help) (regulatory.veevavault.help). These features have already improved regulatory agility. In one case, Moderna moved from Excel spreadsheets to Veeva Vault in under 5 weeks, enabling its small team to handle a **15x** surge in submissions (127 to 2,000+ per year) while giving the entire organization real-time visibility into 1,600+ health-authority queries and past answers ^[10] www.veeva.com ^[11] www.veeva.com).
- **AI opportunities in submission planning:** AI can build on Vault's content-planning capabilities by smarter, data-driven suggestions. For example, machine learning (ML) or generative AI could analyze past submissions and regulatory guidances to recommend which documents to include in a given filing (content plan automation), optimize the filing sequence, or predict realistic timelines. Veeva's own research highlights automating "submission plan refinement" as a high-value use case ^[12] www.veeva.com). AI tools (such as predictive schedulers) could dynamically adjust project timelines by learning typical review and QA durations from historical data. Large language models (LLMs) can assist in drafting parts of the dossier outline or briefing documents, flagging missing sections, and translating regulatory requirements into checklist items. Early pilots show AI drafting of summaries and reports (e.g. Module 2 overviews) can dramatically reduce authors' workload. In one estimate, tailored AI writing reduced CSR drafting time by 40% ^[7] www.mckinsey.com).

- **AI opportunities in health-authority correspondence:** Managing regulatory questions is a persistent bottleneck, but also heavily data-driven. AI techniques are well suited to this domain. For instance, natural language processing (NLP) can **interpret and categorize** incoming queries, automatically flagging urgent issues and routing them to experts. AI platforms – like Freyr's reported "Freya" chatbot – can retrieve past similar questions from a centralized database, eliminating redundant authoring (^[13] www.freyrdigital.com) (^[14] www.freyrdigital.com). Such tools can even generate draft responses: by combining archived Q&A pairs with regulatory guidance, an AI agent can propose an initial answer for human review, cutting response time. Firms are investigating predictive analytics as well: by analyzing patterns of prior submissions, AI could anticipate likely queries for a given product, allowing teams to preemptively address regulators' concerns in the initial filing (^[15] www.freyrdigital.com) (^[16] www.mckinsey.com). In short, AI-augmented query management promises faster, more consistent replies, thereby accelerating approvals (an outcome Freyr notes as a key benefit (^[17] www.freyrdigital.com)).
- **Evidence and case examples:** Numerous industry analyses and early implementations support the potential of AI in RIM. McKinsey's benchmark study (2025) finds only 13–15% of companies have scaled automation beyond basic authoring tasks (^[18] www.mckinsey.com), indicating upside. At the same time, leading firms are piloting advanced solutions. In one published case, a global pharma partnering with SEI (Software Engineering Institute) used a generative AI system to automatically structure legacy submission documents and generate compliant drafts. This resulted in "**automated generation and validation of submission content**", boosting productivity and highlighting missing sections (^[19] www.sei.com). Moderna's Veeva deployment, discussed above, demonstrated the scalability of RIM even without generative AI. And Merck's success with an in-house AI writing tool (cutting CSR drafts by >50% (^[7] www.mckinsey.com)) underscores the practical value of integration.
- **Technological roadmap and implications:** Veeva is actively extending AI into RIM: recent announcements outline "Veeva AI" agents across all Vault applications (^[20] ir.veeva.com) (^[21] ir.veeva.com). For regulatory, planned rollouts (2025–2026) will embed AI agents within Vault RIM workflows (for example, Vault AI eventually will include a "Health Authority Interactions Agent" that automates question-tracking (^[22] www.veeva.com)). These agents are designed to run on secure LLM platforms (Anthropic, Amazon Bedrock) and integrate with Vault data, promising life-sciences-specific intelligence. In parallel, regulatory authorities are crafting policies on AI: the FDA's January 2025 draft guidance stresses "**model credibility**" and validation for AI used in submissions (^[23] www.fda.gov). Thus, implementations must ensure traceability and regulatory compliance. Conceptually, the future regulatory affairs organization could adopt an "AI-augmented human-in-the-loop" model: automated systems draft plans and responses, while experts validate them, ultimately compressing cycle times.

In conclusion, automating submission planning and health-authority correspondence with AI within Veeva Vault RIM holds significant promise. It aligns with current industry trends toward **digitization, data-driven submissions, and agile regulatory processes**. Leading analysts estimate that fully digital end-to-end workflows (with generative AI and real-time data sharing) could enable filing in under six weeks from database lock (^[24] www.mckinsey.com) (^[20] ir.veeva.com). However, realizing this future requires careful attention to data integrity, validation, and change management, ensuring that AI enhancements operate within compliant, validated Vault environments. This report's comprehensive review – drawing on industry benchmarks, Veeva's capabilities, and emerging AI use cases – provides guidance for life-sciences companies planning to harness AI in Vault RIM.

Introduction and Background

Regulatory affairs sits at the nexus of drug development and product approval. By definition, a **regulatory submission** is the complete dossier of data that a sponsor sends to a health authority (FDA, EMA, etc.) to seek approval or clearance of a drug or medical product. Each submission may contain hundreds of documents – from clinical study reports to manufacturing data to labeling – assembled according to standards such as the Common Technical Document (CTD). Historically, these processes were **document-centric and manual**. Sponsors tracked submissions with spreadsheets, prepared PDFs by hand, and relied on attorneys and medical writers for drafting responses to regulators.

While these methods have persistently sufficed, they come with clear drawbacks: slow cycle times, risk of human error, and difficulty sharing information globally. The pharmaceutical industry has long recognized the value of shortening submission timelines. McKinsey estimates that each month shaved off a filing can be worth **tens of millions** in net present value to a company (^[5] www.mckinsey.com). Faster approvals mean patients access therapies sooner and companies maximize patent life. Accordingly, accelerating the submissions process has become a strategic priority, on

par with accelerating clinical development. McKinsey's 2025 analysis reports that leading companies have already **accelerated submissions up to 3x faster** than the 2020 average, with headlines of filing dossiers in 8–12 weeks (down from historical times of several months) ^[6] www.mckinsey.com). Even so, many organizations still struggle with bottlenecks and inconsistency when scaling rapid filings, often due to data silos, complex workflows, or lack of integrated systems ^[25] www.mckinsey.com) ^[26] www.veeva.com).

Regulatory Information Management (RIM) systems emerged over the past two decades to address these challenges. A RIM is a central platform that captures, tracks, and interlinks regulatory data across an organization – including product registrations, submission records, labeling, health authority correspondence, and more. The goal is to replace disconnected processes (email chains, file shares, paper binders) with a unified, auditable repository and workflow engine. According to industry sources, leading biopharma companies began shifting to modern cloud-based RIM as early as the mid-2010s, motivated by evolving regulations like Europe's IDMP (Identification of Medicinal Products) and clinical-trial transparency rules ^[27] www.veeva.com) ^[28] www.veeva.com). For example, one global biopharma implemented a new cloud RIM in 2016 to centralize **all** regulatory content across clinical, quality, and submissions domains, after realizing that “fragmented data and documentation... had slowed its ability to respond to global regulatory and business changes” ^[26] www.veeva.com) ^[29] www.veeva.com). With RIM, that company was able to automate and integrate processes and update their system multiple times per year – a huge leap from the one-change-per-year cycle under legacy systems ^[30] www.veeva.com).

Veeva Systems' Vault RIM suite has become one of the industry's most widely adopted RIM solutions (Veeva reports usage by **15 of the top 20** pharmaceutical companies ^[4] www.veeva.com), and over 400 global firms ^[31] www.veeva.com). Vault RIM includes modules for **Submissions** (for document planning/assembly), **Interactions** (for tracking communications with regulators), **Registrations** (for product registration data), and more. Vault's cloud platform offers features such as content co-authoring, version control, and real-time dashboards, helping regulatory teams rapidly build submission dossiers and monitor their progress. For example, Vault's **Submissions** application allows users to “plan, author, review, and approve regulatory documents” in a structured way ^[32] www.veeva.com). It provides content planning capabilities: users can “build a submission outline and automatically match documents to the outline” ^[32] www.veeva.com). Vault Submissions also supports report-level content plans and publishing automation (e.g. to FDA's ESG), enabling continuous electronic publishing ^[33] www.veeva.com) (regulatory.vaultbasics.help). On the interactions side, Vault RIM offers a **Health Authority Interactions** (HAI) panel that can scan documents (like letters or meeting minutes) for questions or commitments. When a user highlights a question in a correspondence document, Vault automatically extracts it as a **Health Authority Question** (HAQ) record, and links it to the relevant submission and product (regulatory.veevavault.help). The system then guides users through issuing a **Health Authority Response** (HAR) document: a template-based reply can be generated and attached to the related HAQs (regulatory.veevavault.help).

These RIM features have delivered measurable efficiency gains. For instance, a published customer story describes how Moderna's small regulatory team used Vault RIM to replace spreadsheet-based tracking of COVID vaccine Q&A. Prior to implementing Vault, Moderna regulators “often spent valuable time addressing [similar questions] separately” for different submissions, with the answers stored in disjointed spreadsheets ^[34] www.veeva.com). After going live with Vault in under five weeks, Moderna personnel could easily search over **1,600 queries and answers** in the Vault, quickly identifying duplicate questions and re-using prior responses. This unified database not only **streamlined response drafting** but also improved consistency across global answers ^[11] www.veeva.com). The team even structured assignments so that HAQs could be routed to the right experts with due dates, allowing them to handle a 15-fold surge in submission volume (from 127 to over 2,000 annually) without slowed review cycles ^[35] www.veeva.com). In short, RIM has helped break down silos and enable scale in regulatory affairs.

Nevertheless, many RIM functionalities remain largely **rule-based and manual**. For example, Vault's content plans must still be created from static templates, and matching documents to the plan is a user-driven task. Responding to HAQs, while supported by guided workflows, still relies on human authors to draft each answer based on templates. Routine steps like version control, formatting, and compliance checks, though automated in parts, still consume significant human time. This status quo presents an opportunity: adding **AI and advanced automation** into RIM workflows can further

compress timelines and reduce human workload. The rest of this report explores how such AI integration could power submission planning and HA correspondence within Veeva Vault RIM, drawing on academic studies, industry analyses (like McKinsey [3] and regulatory technology vendors (^[36] www.freyrdigital.com) (^[37] app.submitly.ai)), and real-world implementations. We analyze current capabilities, potential AI use cases, and strategic considerations, including data, compliance, and organizational impact. The goal is to provide a comprehensive, evidence-based picture of “AI for Vault RIM,” focusing on the critical tasks of submission planning and regulatory correspondence.

Regulatory Submission Planning in Vault RIM

The Submission Planning Process

Regulatory submission planning is the process of defining what documents and data need to be sent to each health authority, and when. A major submission (e.g. an NDA, MAA, or variation dossier) typically consists of hundreds of files, organized into modules (e.g. eCTD Modules 1–5). Planning involves deciding the content checklist, scheduling preparation of each document, and coordinating across functions (clinical, CMC, safety, labeling, etc.). Planners must translate high-level regulatory strategy into concrete deliverables – for example, mapping which clinical study reports, statistical analyses, or manufacturing reports will support each indication. The plan must also adapt to each country’s requirements: some regions might need additional local language labels or country-specific forms, others might waive certain modules. Historically, companies often handled this with complex spreadsheets or documents listing the submission contents and status, which risked version skew and duplications.

Veeva Vault RIM’s Submissions module is designed to replace that paradigm with a structured **Content Plan**. A content plan is a hierarchical outline (essentially, a “Table of Contents”) representing all expected documents for a submission ([regulatory.vaultbasics.help](#)). Veeva administrators set up a template for the content plan (often based on a common dossier format). When a submission record is created in Vault (e.g. an NDA targeting FDA), a user can click *Create Content Plan* to generate the outline automatically ([regulatory.vaultbasics.help](#)). Vault then populates the plan with placeholder entries (binders, sections, leaf documents) according to the template. Each entry can be linked to an actual file in the Vault library; as authors draft documents, they match each completed document object to the corresponding plan item. The content plan can be viewed in a spreadsheet-like panel, showing which items are complete, in progress, or pending. This automated outline creation reportedly “auto-generates a table of contents for major regulatory submissions” and lets submission managers **track real-time status** of each document (^[9] www.veeva.com) ([regulatory.vaultbasics.help](#)). In practice, this means a submission manager always has a single source of truth for which modules are covered and which files are outstanding.

Beyond basic listing, Vault provides some automation for content plans. For example, if a binder in the plan is linked to a Vault folder or library, new documents placed in that folder can be auto-matched to the plan. Vault can also generate initial document objects for each plan item (using form metadata), saving manual data entry ([regulatory.vaultbasics.help](#)). After database lock, Vault’s **Submissions Publishing** feature can take the finalized content plan and render the documents into an electronic format (eCTD or NeeS) for submission. Thanks to this workflow, Vault RIM has been shown to eliminate redundant uploads and reduce manual entry. One regulatory leader said, “For most system users, [Vault] will make their lives easier because they don’t have to upload duplicate documents or manage manual data entry” (^[38] www.veeva.com).

Despite these advances, the submission planning process in Vault still relies on human judgment. The initial content plan template must be configured in advance by RIM administrators. When new products or indications arise that don’t match existing templates, teams often update the template. The system doesn’t itself decide *what* categories are needed for a brand-new dossier; it follows rules given by users. Planners must still identify which documents should be included (for example, selecting which of dozens of clinical reports to include for a supplemental filing). Timeline scheduling is outside

the scope of content plans: Vault tracks content readiness status but does not inherently model *when* documents will be ready. Submission managers commonly use external project plans or checklists to coordinate tasks.

In summary, Vault RIM's submission planning tools (content plans, binders, dashboards) provide strong structure and visibility to the planning process (^[9] www.veeva.com) ([regulatory.vaultbasics.help](#)). They ensure teams work from a shared outline and reduce lost files/errors. However, tasks such as choosing which items to include, prioritizing critical documents, and adapting plans dynamically remain largely manual. This leaves room for AI-driven tools to augment planning in several ways – for example, by analyzing historical submissions to suggest the optimal set of contents, or by predicting resource requirements and deadlines. We explore those AI opportunities in the next section.

AI Opportunities in Submission Planning

AI and automation offer multiple avenues to enhance submission planning beyond what static templates can do. Broadly, AI can assist with **content identification, schedule optimization, and knowledge synthesis**:

- **Content Identification and Generation:** By mining prior submissions and regulatory databases, an AI model can help determine which components are needed in a new dossier. For example, if a company has previously filed a similar product, an AI agent could compare product profiles and regulatory objectives, and suggest a draft content plan by analogy. Generative AI (large language models) can even propose document titles or section descriptions based on the submission's context. Early prototypes use-case include drafting initial versions of overview documents. In practice, Veeva RIM already provides a starting outline, but AI could automatically **tailor** that outline to the specific product. For instance, Vault's "Content Plan Template" could be dynamically adjusted: an AI agent could insert or remove sections (such as pediatric studies, post-market commitments, or regional annexes) by checking global regulatory requirements for the product's attributes. Indeed, internal Veeva research identifies "submission plan refinement" as an attractive automation use case (^[12] www.veeva.com). In time, we may see AI-driven content planners that continuously learn from each filing, spotting opportunities to reuse approved content or flag obsolete items.
- **Timeline Prediction and Management:** Submissions are not just static outlines; they follow tight calendars (e.g. completing each module by database lock). AI can enhance scheduling by analyzing project metrics and historical timelines. For example, machine learning models could ingest past submission records (in Vault or external project tools) to forecast realistic document completion times. If a draft clinical study report historically took 30 days to finalize, an AI scheduler could use that estimate to set deadlines and flag at-risk tasks. Predictive analytics could also simulate parallel workflows: for instance, AI could recommend "parallel authoring" of certain sections (as advocated by lean processes (^[39] www.mckinsey.com)). Though specific AI timeline calculators are still emergent, consultants already highlight the prospect of AI timeline tools ("calculate your timeline edge") for regulatory affairs (^[40] [internal.ai](#)). A data-driven submission planner would alert managers early if a filing is slipping and could propose mitigation steps (e.g. overlap reviews or shift low-priority documents to subsequent submissions).
- **Knowledge Extraction and Insights:** Regulatory submissions draw on mountains of guidance (agency guidances, ICH documents, local regulations) and prior intelligence. NLP and knowledge-graph techniques can help synthesize this knowledge. For example, a company's proprietary regulatory intelligence repository (HA guidelines, minutes of prior meetings, internal Q&A) could be fed into an AI summarizer. The AI could then distill high-level recommendations or regulatory "requirements" that should influence the submission plan. Concretely, it might flag that a certain formulation change triggers extra chemistry reports, or that a new regulatory pathway in a country requires a justification document. Veeva's own blog notes that GenAI is being explored to "summarize large volumes of unstructured HA guidance documents and internal local intelligence" (^[41] www.veeva.com). Embedding such analysis within Vault could automate tasks like tagging relevant guidance to a product's regulatory strategy.
- **Document Drafting Support:** While writing is often considered a separate step, AI can start the moment planning ends. Large language models can generate draft document templates and boilerplate text for modules commonly used in submissions (e.g. Module 2 summaries, cover letters). Veeva's partner Merck reported that an AI writing tool cut first-draft CSR time from 180 to 80 hours (^[7] www.mckinsey.com). In the planning phase, AI might auto-fill parts of the content plan with key document metadata (e.g. proposed titles, responsible authors). The AI could also check consistency between different documents: for instance, ensuring that key statements in the Clinical Overview align with results presented in individual study reports. This type of cross-fingerprinting could catch missing references early, effectively "validating" the submission outline.

These AI enhancements all build on Vault's data model. Because Vault stores structured fields (e.g. product details, regulatory objectives, submission types), an AI built into the platform can condition on real data – not just free text. Veeva's proposed AI agents will be "application-specific" with direct access to Vault data (^[20] ir.veeva.com) (^[21]

ir.veeva.com). For example, a Regulatory AI agent could use the Vault **Regulatory Objective** and **Submission** records to tailor its suggestions. In practice, a submission planner might interact with an AI “assistant” through chat or form-based prompts: “Generate a content plan for Product X’s new indication in Region Y,” and the system would return an initial plan outline for user review. Such interaction could live natively in Vault (for instance, as an “AI Shortcut” in the user interface ([42] ir.veeva.com)).

The expected **benefits** of AI-augmented planning are clear. Automating the outline creation and document matching can reduce the manual workload and rework of draft plans. One industry analyst notes that AI in submissions can “ensure consistency in responses across different regulatory markets” and “minimize redundant work by suggesting pre-validated responses” (for HA queries) ([14] www.freyrdigital.com) – a principle that also applies to planning. Preliminary pilots (both internal and from vendors) suggest that companies can cut critical path by front-loading and automating document assembly ([43] www.mckinsey.com) ([17] www.freyrdigital.com). In numerical terms, a McKinsey study estimates that companies using data-centric automated workflows (with structured authoring and AI) are able to file in **under 8 weeks after database lock**, compared to the historical 6–12 months ([6] www.mckinsey.com) ([44] www.mckinsey.com). While every organization’s baseline differs, even partial automation (rule-based plus AI assistance) can shave weeks or months off filing timelines. In the next section, we will quantify some of these gains with data and case examples, and begin to outline transformation roadmaps.

Table 1: AI/Automation in Regulatory Workflow Tasks

Task	Traditional Process	Vault (Pre-AI) Process	AI/Automation Opportunities & Benefits
Submission Content Planning	Teams build outlines manually (often in Excel or word) based on general CTD templates. Determining which modules/docs apply requires cross-functional coordination. Plans updated via email/sheets, error-prone.	Vault auto-generates a hierarchical content plan (tablename of contents) from a template (regulatory.veevavault.help). Users assign actual documents to plan items; dashboards track status of each item. No built-in timeline modeling.	<i>AI Content Advisor</i> : uses past submission data to suggest missing sections or country-specific modules, generating an adaptive content plan. <i>Generative Overview Drafting</i> : LLMs draft protocol synopses or summaries based on available data. <i>Benefit</i> : Reduces manual outline edits and omissions, accelerates plan completion.
Document Assembly & Authoring	Authors draft each document in Word; formatting and PDF conversion manual. Version control via filenames. Consolidation steps (bookmarking, hyperlinking) done manually in publishing tools.	Vault provides real-time collaborative authoring, version histories, and auto-matches docs to plan. The Vault Submissions Publishing tool automatically compiles and validates documents (generating eCTD packages).	<i>Intelligent Drafting</i> : AI suggests boilerplate text or sections for common document types (e.g. protocol summary, eCTD introduction), based on stored data. <i>Automated QA Checks</i> : ML models flag inconsistencies or missing citations before publishing. <i>Benefit</i> : Cuts authoring/review cycles (e.g. Merck’s generative AI cut CSR writing time ~56% ([7] www.mckinsey.com)).
Submission Timeline Management	Project managers create external timelines/Gantt charts; often isolated from Vault. Milestones tracked in project tools (e.g. weeks from DUE). Delays not visible in Vault itself.	Vault tracks submission readiness status (documents approved/pending) but does not simulate schedule. System monitors content plan completeness.	<i>Predictive Scheduling</i> : ML forecasts document completion times and signals bottlenecks. <i>Scenario Simulation</i> : AI recommends fast-track strategies (e.g. parallel review, outsourcing) via what-if analysis. <i>Benefit</i> : More reliable DBL-to-filing forecasts; proactive mitigation of delay risks (thus enabling goals like 8–10 week filing ([44] www.mckinsey.com)).
Health Authority Questions (HAQs)	Regulatory queries logged in spreadsheets or siloed email threads. Reviewers manually search past Q&A or guidance for similar issues. Drafting answers often begins from scratch.	Vault’s Health Authority Interactions panel lets users extract HAQs from correspondence and link them to submissions (regulatory.veevavault.help). HAQ records track related app, submission, and response. Vault can auto-create a response record and template document (HA Response) when needed (regulatory.veevavault.help).	<i>NLP Query Triage</i> : Automatically categorize and prioritize incoming queries by urgency/topic ([36] www.freyrdigital.com). <i>Intelligent Retrieval</i> : AI suggests relevant past questions/answers as references ([13] www.freyrdigital.com). <i>Draft Generation</i> : LLMs generate initial HAQ responses using previous answers and guidelines ([16] www.mckinsey.com). <i>Benefit</i> : Faster HAQ turnaround (AI can “cut HAQ response time by automating classification, retrieval of past responses, and drafting” ([17] www.freyrdigital.com)), improved consistency across similar queries.
Regulatory Intelligence (Guidance)	Regulatory affairs teams manually search FDA/EMA guidelines for clues. Each question requires re-reading lengthy documents.	Vault stores attachments of major regulations/communique, but analysis remains manual. No AI indexing of these docs by Vault’s core.	<i>Automated Summarization</i> : AI summarizes key guidance relevant to the product profile. <i>Chatbot Advisor</i> : Interactive agent answers regulatory strategy questions using company knowledge base. <i>Benefit</i> : Quicker insights on regulatory requirements, reducing strategist time from days to hours ([41] www.veeva.com).

Note: Table 1 summarizes representative areas where AI or advanced automation can add significant value to the submission process. The cited benefits have been observed in industry studies and pilot projects ([7] www.mckinsey.com) ([17] www.freyrdigital.com).

Health Authority Correspondence in Vault RIM

The HA Correspondence Workflow in Vault

A vital component of regulatory operations is **health-authority correspondence**: the questions, comments, and commitments exchanged with agencies during the review of a submission. In practice, a submission often triggers multiple rounds of queries from regulators ("Health Authority Questions" or HAQs) and requires formal written responses (respectively, "Health Authority Responses" or HARs). Additionally, sponsors prepare meeting packages and transcripts when engaging with authorities. Tracking these interactions meticulously is crucial for compliance and to avoid missing follow-up requests.

Veeva Vault RIM includes a built-in **Health Authority Interactions** feature to manage this process. Conceptually, each correspondence document (letters, meeting minutes, emails) can be scanned to extract the specific Q&A interactions. As the Vault Help describes, a user highlights a question sentence in a document, and Vault automatically creates a linked HAQ record ([regulatory.veevavault.help](#)). The system annotates the source document and suggests linking that HAQ to the relevant submission, application, and product in the Vault. Consequently, all HAQs are stored centrally as structured records, rather than buried in PDFs. Vault likewise captures **Commitments** (promises to regulators) if text is highlighted as such. Each HAQ record has metadata fields (e.g. priority, owner, due date) and is an object that can be queried. Users can filter and review all HAQs across products in a consolidated view.

Once HAQs are in Vault, the next step is to prepare official **responses**. Vault automates part of this flow via the *Initiate Response* function ([regulatory.veevavault.help](#)). On an HAQ, the user can click "Initiate Response," launching a guided multi-step tool. Vault will: create a new HAR record, link it to the application and submission, generate a formatted Word document from a response template, and (optionally) create a new submission (for the response) with its own content plan ([regulatory.veevavault.help](#)). Users select which HAQs to associate with this response, order them, and Vault then produces an empty response letter with placeholders for answers. The HAQs are automatically linked to the HAR record and the new submission. Vault's logic ensures that all HAQs in that view share the same application, so responses remain consistent. Once created, the response document can be co-authored by the relevant experts. After finalizing the response, it is routed (often through QA and legal approvals) and eventually submitted to the authority. Vault tracks that a submission with the "HA Response" dossier is in progress and archives the outgoing letter against the original submission's record.

These features have significantly improved visibility and organization. Regulators' questions no longer live in scattered inboxes; they are captured as database records and linked to all relevant artifacts. Users can easily see, for example, all open queries by product or country, how many have been answered, and where delays exist. Workflow assignments and due dates on HAQs ensure accountability. As Moderna's CFR instance illustrated, this unified approach "enhanced visibility into past responses and current queries" and "streamlined responses" so their team could keep pace with massive submission volume (^[10] [www.veeva.com](#)) (^[35] [www.veeva.com](#)).

Nonetheless, despite Vault's keyword highlighting and structured HAQ objects, drafting the actual answers remains manual. The system helps organize which questions need replies and generates a document shell, but each answer still requires research, writing, and approval by experts. In large filings, multiple regulatory queries might require input from different departments (e.g., clinical, safety, manufacturing). Coordinating those answers and ensuring consistency is time-consuming. Additionally, Vault only extracts questions marked by a user – it does not autonomously parse incoming letters. Non-text elements (scanned PDFs, meeting transcripts) require manual review. And critically, Vault's HA features do not inherently predict or preempt queries; they react to questions after the fact.

These pain points strongly motivate intelligent automation. Given the volume of text and similarity of inquiries, AI can potentially shoulder much of the heavy lifting. The next section explores how AI technologies can transform health-authority correspondence within Vault RIM.

AI-Powered Health Authority Correspondence

AI offers several promising enhancements to the HA query/response process:

- **Automated Query Intake and Triage:** Instead of waiting for a user to highlight questions, AI (especially NLP and OCR) can scan incoming correspondence to **automatically detect and categorize queries**. For example, an email or PDF letter could be parsed by an AI agent that identifies question sentences and creates HAQ records without manual highlighting. Freyr notes that modern RegTech can “automate the end-to-end query management process,” including categorizing and prioritizing queries by urgency (^[36] www.freyrdigital.com). In practice, an AI could flag high-priority questions (e.g., clinical safety concerns) and route them immediately. If multiple questions share a theme, the system could batch them to one sub-team. The goal is to eliminate the initial manual step of recording HAQs, reducing delays in acknowledging regulator requests.
- **NLP and Semantic Search for Similar Queries:** Perhaps the most immediate impact is leveraging past Q&A to speed drafting new responses. Vault RIM’s HAQs, combined with older HAQs and responses, form a rich knowledge base. AI can index this content semantically: by running embeddings or NLP on question/answer pairs, a system can identify which historical queries are most similar to a new one. For instance, if an EMA reviewer asks about drug interactions, the AI can retrieve past HAQs on interactions in other dossiers, along with how they were answered. Freyr’s Freya chatbot is an example: it lets regulatory teams “access real-time regulatory insights, automate data retrieval, and assist in crafting accurate responses to HA queries” by tapping into such repositories (^[45] www.freyrdigital.com). The net effect is to prevent duplication of effort: if a question has been effectively answered before, authors can reuse or adapt that text. Moderna’s implementation already gave users a filtered view of 1,600 similar Q&A records (^[11] www.veeva.com); AI can automate the filtering process and surface the best matches instantly.
- **Automated Drafting of Responses:** Beyond retrieval, generative AI can propose draft answers. Given a question and the relevant product data, an LLM (trained or fine-tuned on regulatory writing) could compose a preliminary response. It might draw from the company’s data, published literature, and internal guidelines to form an answer outline. Of course, due to compliance requirements, such AI-drafted text would still undergo expert review. But early pilots show high potential: McKinsey suggests GenAI could be used to “generate HAQ responses, including for situations when multiple health authorities submit queries simultaneously” (^[16] www.mckinsey.com). This capability would be especially powerful during global filings when, for example, an FDA and EMA query the same issue. An AI could draft one coherent answer adapted to each agency’s format, saving duplicated writing. Even on a smaller scale, a feasibility study by SEI demonstrated that a generative AI workflow “automatically [generated] first-draft submission recommendations grounded in historical context and current standards” (^[19] www.sei.com), which presumably would include answering anticipated questions. In practice, we might see features where, on clicking “Draft Response”, Vault fires an AI that outputs a word document addressing the selected HAQs, subject to human edits.
- **Summary and Predictive Analytics:** AI can look ahead at the submission package and regulatory history to *predict* the kind of questions likely to arise (for example, by running risk models on the dossier content). This is inherently suggestive: if such predictions are reliable, teams can proactively include clarifications in the original submission, reducing future queries. Freyr 2025 commentary highlights this as a benefit of AI: by analyzing past interactions, AI can “predict potential queries that may arise during submission review... allowing teams to preemptively address common Health Authority concerns in initial submissions” (^[15] www.freyrdigital.com). In Vault, such functionality could manifest as a “Query Predictor” dashboard: after the submission plan is complete, the system flags certain HAQs (with estimated probability/confidence) so document authors know to double-check those areas.
- **Knowledge Management and Consistency:** With AI-powered knowledge bases, consistency across correspondence is easier. For example, once a regulatory preference (e.g. “EMA prefers X wording”) is learned and stored, the AI can ensure all new communications adhere to it. AI can also maintain a living glossary of terms and commitments per agency. The aforementioned Freya model also “ensures consistency in responses across different regulatory markets” and “minimize [s] redundant work” by suggesting pre-validated templates (^[46] www.freyrdigital.com). Internally, this could mean Vault RIM’s HA module uses machine learning to cluster similar questions together, alerting reviewers if two answers differ in content and should be harmonized.

Table 2 (below) summarizes some of these AI-based enhancements and their effects on approval speed and workload. Notably, reducing query response time has outsized value: Freyr lists “**Faster Response Time**” as the #1 benefit of AI in HA query management (^[17] www.freyrdigital.com), noting that automating classification and draft writing “leads to quicker query resolution and faster market approvals” (improving first-pass success rates). Indeed, McKinsey estimates that accelerating each submission by even a few weeks has multimillion-dollar payoffs (^[5] www.mckinsey.com), much of which can come from cutting query cycles.

Table 2: AI-Powered Health Authority Query Management

HA Task	Current (Vault RIM)	AI-Enhanced Process	Impact on Review Cycle
Query Detection	Manual input of HAQs by highlighting docs (regulatory.veevavault.help); possible delays if questions are overlooked.	Automated parsing of incoming letters/emails to detect questions and populate HAQ records (OCR+NLP). Auto-tag by category (efficacy, safety, etc.).	Faster initiation of QA workflow; eliminates lost/missed queries.
Prioritization & Routing	User-assigned priorities and manual assignment. Reliant on project managers to triage.	AI classifies queries by urgency/complexity (^[36] www.freyrdigital.com), dynamically prioritizes and assigns to subject experts.	Ensures critical questions get addressed first, reducing hold-ups.
Response Drafting	Authors draft answers in Word based on SOP templates. Manual review of past responses for reference.	AI retrieves similar past Q&As from Vault (^[13] www.freyrdigital.com); generates a first-draft HAR response (^[16] www.mckinsey.com). (subject to editing)	Cuts author time; reduces errors; can handle simultaneous global queries consistently.
Knowledge Retrieval	Search by keyword or manual review; unstructured guidance review.	AI-powered search: semantic/embedding search over regulatory intelligence and past HAQs. Provides relevant guidelines and prior answers instantly (^[46] www.freyrdigital.com).	Improves accuracy and completeness of answers; avoids redundant research.
Predictive Analytics	None; reactive process.	AI models estimate which HAQs are likely (e.g. identify risky sections of dossier).	Teams address key issues up-front, improving first-cycle approval rates (^[15] www.freyrdigital.com).
Quality & Consistency	Dependent on manual QC; risk of inconsistencies if multiple authors.	AI-driven consistency checks (flag contradictions, terminology mismatches). Use of centralized "knowledge" ensures uniformity (^[14] www.freyrdigital.com).	Reduces need for multiple review rounds; strengthens compliance.

Sources: Industry analyses and case studies indicate that AI can automate much of the HA query workflow (^[36] [www.freyrdigital.com](#)) (^[16] [www.mckinsey.com](#)) and that such automation yields measurable speed and quality gains (^[17] [www.freyrdigital.com](#)) (^[5] [www.mckinsey.com](#)).

Integrating AI with Veeva Vault RIM

Successful AI implementation requires integration with existing systems. Fortunately, Veeva is proactively preparing Vault for AI augmentation. In April 2025, Veeva announced a strategic initiative called “**Veeva AI**”, aimed at embedding intelligent agents and shortcuts across the Vault Platform (^[20] [ir.veeva.com](#)). Veeva AI is explicitly *LLM-agnostic*, allowing customers to use embedded models or plug in their own via cloud AI services (^[47] [ir.veeva.com](#)). By late 2025, early AI features (e.g. in Vault CRM) became available, and a timeline for regulatory was set: *Deep, industry-specific AI agents* targeting regulatory workflows are slated for late 2026 (^[21] [ir.veeva.com](#)). Specifically, Veeva’s press release (Oct 2025) indicates an **August 2026** agent rollout for regulatory/medical applications (^[21] [ir.veeva.com](#)). These agents will be able to “see” the content of Vault – including Regulatory Objectives, HAQs, submission plans, and attachments – and apply business-templated LLM prompts.

Veeva also plans user-extensible AI: customers can configure the delivered agents or build custom “AI Shortcuts” in Vault (^[48] [ir.veeva.com](#)) (^[49] [ir.veeva.com](#)). For example, a regulatory user could define a shortcut like “Draft Answer for Selected HAQs,” which invokes an AI agent only on the selected HAQ context. Because Veeva’s AI lanes are usage-metered, companies can start small (e.g. enabling one agent for beta) and scale later. Under the hood, Veeva’s agents will run on secure LLM servers (Anthropic/Bedrock/Microsoft AI Foundry) and communicate via Vault’s API (^[21] [ir.veeva.com](#)), keeping data sequestered per customer. The system also promises prompts and safeguards tuned to regulatory language (e.g. avoiding hallucinations in a low-risk way).

From a technical standpoint, deploying AI in a validated regulatory system involves several considerations:

- **Data Privacy and Security:** Vault stores highly confidential content. Veeva AI architecture uses private models and encrypted connections to prevent baseline model leaks (^[47] [ir.veeva.com](#)). Companies may choose to use on-premise or private cloud models for sensitive data.

- **Validation and Compliance:** Any AI feature must be validated under GxP guidelines. Outputs from AI (answers, summaries) will likely need human verification, which means creating logs showing how the AI was invoked and that results were reviewed. Veeva's AI platform will preserve audit trails by design, but sponsors will want to document that, for example, every AI-drafted response was vetted by a qualified medical writer.
- **Model Training / Fine-tuning:** Regulatory writing is highly specialized. Companies should consider fine-tuning LLMs on their own past dossiers and HAQ archives, or restricting generation to approved safety documents only (e.g. using an LLM PHR based on internal cits). Veeva's partner ecosystem can help build such fine-tuned models.
- **Change Management and Training:** RIM users will need training to trust and properly use AI tools. Veeva is likely to provide in-app tutorials (as it has with new platform features) to guide users through the first AI interactions. Regulatory teams must also establish procedures for editing and approving AI suggestions, and possibly new roles (e.g. an "AI reviewer").

It is also worth noting that AI integration in Vault is designed to complement (not replace) human expertise. Veeva emphasizes that AI Agents will operate "**alongside a human in the loop**" ([49] ir.veeva.com). In McKinsey's framework, industry leaders envision an "agentic AI" system that challenges content while experts make final decisions ([16] www.mckinsey.com). Practically, this means if an AI generates a dossier section or query answer, a knowledgeable RA reviewer is still responsible for the final text. This hybrid approach mitigates the risk of AI hallucinations by leveraging the reviewer's domain knowledge ([50] pmc.ncbi.nlm.nih.gov).

Throughout all these developments, the underlying Vault platform remains the single source of truth. AI does not operate in isolation; any plan it creates, any HA response it drafts, flows back into Vault as a document or record. For example, if an AI recommends adding a certain document to the content plan, that addition would manifest as a change in the Vault content plan object, subject to change control. Similarly, AI-suggested text in an HA response can be directly edited in the Vault Word editor, with versions tracked. This tight integration will ensure auditability and regulatory compliance: every change, whether human or AI-assisted, is recorded in Vault's logs.

Data Analysis and Evidence-Based Insights

To quantify the impact of AI and automation in RIM, we draw on industry benchmarks and published results:

- **Timeline Reduction Benchmarks:** McKinsey's recent analysis provides concrete figures. In 2020, the average DBL (Database Lock) to filing time for submissions was on the order of **24–30 weeks** across the industry. Top performers have already slashed this by over half. Some leaders now deliver filings **8–12 weeks** post-DBL ([6] www.mckinsey.com). McKinsey outlines a three-horizon transformation path (Exhibit below) where:

- *Horizon 1 (Process revamp)* yields DBL-to-file in ~12–14 weeks ([44] www.mckinsey.com).
- *Horizon 2 (Full operating model redesign)* targets ~8–10 weeks ([44] www.mckinsey.com).
- *Horizon 3 (Full digital and AI adoption)* aspires to *under 6 weeks* ([24] www.mckinsey.com).

The table below (Table 3) summarizes these benchmarks. It illustrates that with advanced automation (including AI content generation and continuous publishing), companies aim for filing in a few weeks rather than many. These benchmarks are consistent with Veeva customer reports of accelerated publishing cycles. As one Veeva customer put it: "We have this strong foundation to bring all of our data, content, [and] processes... and have it fully integrated" ([51] www.veeva.com), enabling much faster execution.

Table 3: DBL-to-File Timelines with RIM Transformation (Weeks)[3] [15]

Transformation Stage	DBL-to-Filing Time (Weeks)	Notes
Baseline (2020 industry avg.)	~24–30	Typical pre-automation timeline.
Horizon 1 (Process Revamp)	~12–14	Cordoned "zero-based" processes (e.g. aggressive front-loading, lean writing) ([44] www.mckinsey.com).

Transformation Stage	DBL-to-Filing Time (Weeks)	Notes
Horizon 2 (Operating Model Redesign)	~8-10	Full submissions process re-engineered, including partial digital integration and GenAI pilots (^[44] www.mckinsey.com).
Horizon 3 (Fully Digital/AI)	<6	Ambitious goal with real-time data and submission (e.g. FDA's Real-Time Oncology Review) (^[24] www.mckinsey.com).

Sources: McKinsey & Co. (2025) analysis of industry submission timelines (^[44] www.mckinsey.com) (^[24] www.mckinsey.com) and public announcements on AI initiatives (^[20] ir.veeva.com). Note that actual times vary by asset, company, and indication; the above illustrates the potential compression of cycles as firms adopt Vault/AI strategies.

- **Capacity Gains from Automation:** Data on manual effort highlights the current gap. Industry surveys show that, beyond core authoring tools, only about **10–15%** of companies have scaled automation in areas like table generation, quality checks, or submission assembly (^[18] www.mckinsey.com). Many organizations still spend weeks on formatting, indexing, and checking each dossier. Freyr quantifies the problem: it reports that “manual bookmarking, hyperlinking, checking PDF properties... leads to significant time consumption,” causing submission delays (^[2] www.freyrsolutions.com) (^[52] www.freyrsolutions.com). By contrast, even early RPA scripts have shown that automating tasks like automatic bookmarking or batch PDF generation can cut days of pole work. This suggests a large upside from further automation, once deployed across scaled pipelines.

Regarding health-authority queries, anecdotal metrics illustrate AI's promise. Moderna's example highlighted that one team went from 127 to 2,000+ submissions per year, yet “kept approvals on track” by using Vault RIM to manage 1,600 queries (^[35] www.veeva.com). Although this was not explicitly AI-driven, it establishes a baseline: the new process handled 15x submission volume without explosion of workload. If AI further reduces hours per query, this productivity could multiply. Freyr's analysis emphasizes that AI can **reduce turnarounds** by automating query workflows, projecting that top-line approvals speed up accordingly (^[17] www.freyrdigital.com).

- **Case Study: Moderna (VAa** RIM Implementation)**** – Moderna's experience is illustrative. Before Vault, their small team tracked queries in disconnected spreadsheets, causing duplicative effort (e.g. 3 health agencies asking similar questions (^[34] www.veeva.com)). Deploying Vault RIM in <5 weeks gave them a “more scalable HA query management process” (^[10] www.veeva.com). Quantitatively, they now filter and reuse responses among >1,600 QA entries (^[11] www.veeva.com), enabling them to deal with a **15x surge** in amendments without major delays (^[35] www.veeva.com). Moderna's story shows that even without AI, a centralized RIM data model can dramatically improve efficiency. Layering AI on top of this—for instance, by auto-filling answer drafts from those 1,600 records—could yield another order of magnitude savings.
- **Case Study: Generative AI Pilot (SEI & Pharma Client)** – A 2025 case study by SEI (Software Engineering Institute) experimented with GenAI for submission drafting (^[19] www.sei.com). A global pharma partnered with SEI to “accelerate regulatory submission workflows” using AI. The solution involved structuring historical submissions and tuning LLM prompts with regulatory context. Results were compelling: cutting manual effort, generating prioritized draft content, and flagging missing criteria in submissions (^[19] www.sei.com). While exact numeric gains were not publicized, the project's reported outcomes (“boosted workforce productivity” and automated validation) align with building-block analyses in [3] – reinforcing that combining data and AI can streamline reviews.
- **Industry Surveys and Expert Opinions:** Besides corporate studies, independent experts highlight these trends. A 2023 Veeva blog noted that leading biopharma are investing heavily in GenAI “proof of concepts” for regulatory tasks (^[53] www.veeva.com). They urge caution (noting hallucination risk), but the same analysis detailed that GenAI has been tested for things like dossier authoring and query response generation (^[54] www.veeva.com) (^[55] www.veeva.com). Separately, Freyr's thought leadership (February 2025) enumerates the tangible benefits of AI in query management – including “improved accuracy and compliance” and “resource optimization” (^[17] www.freyrdigital.com). These sources collectively provide converging evidence: companies that systematically apply structured content and AI-enabled workflows will see measurable acceleration of submissions (and avoidance of recall cycles) compared to those relying on manual documents (^[56] www.mckinsey.com) (^[17] www.freyrdigital.com).

Data Visualization: Publication Timelines (Figure 1)

! [Figure 1: Example submission pacing across transformation horizons (^[44] www.mckinsey.com) (^[24] www.mckinsey.com)]
 (_IMAGE_DATA_HERE)

Figure 1: Illustrative firm benchmarks for Database-Lock-to-Filing time under different submission excellence transformations. Current industry averages linger around 20–24 weeks (Bar 1), while leaders achieve 8–12 weeks. The ultimate target (Horizon 3) is under 6 weeks (^[44] www.mckinsey.com) (^[24] www.mckinsey.com). (Data adapted from McKinsey 2025.)

Note: This figure is schematic, intended to visualize the scale of improvement. Actual results vary by company and product.

Case Studies and Real-World Examples

To ground the above analyses, we review several real-world examples of RIM automation (with and without AI) in life sciences companies:

- **Moderna (Health Authority Queries)** – *Scenario:* Rapid-fire COVID-19 submissions outpaced Moderna's legacy query process. *Intervention:* Implement Vault RIM (Submissions + Interactions) in a 5-week deployment. *Outcome:* Unified HAQ tracking allowed immediate identification of duplicate queries. Queries are now assigned to teams with alerts. Moderna's senior RA lead noted that VAULT "streamlined responses to keep pace with a 15x increase in submissions." (^[10] www.veeva.com) (^[57] www.veeva.com). They report the ability to filter >1,600 Qs/As across products, greatly increasing consistency (^[11] www.veeva.com). *Implications:* This case demonstrates the value of consolidating query data. (AI could further enhance by auto-filling answers from the vault of past responses, but even without AI, productivity soared.)
- **SEI/Global Pharma (Submission Content Drafting)** – *Scenario:* A large pharma's new drug submission was delayed by bottlenecks in authoring and reviewing module documents. *Intervention:* SEI collaboration using generative AI to convert historical dossiers into structured data and generate draft submission text. Heavily bespoke prompt engineering and iterative validation were applied. *Outcome:* The AI pipeline "boosted workforce productivity by reducing manual effort" and "generated first-draft submission recommendations grounded in historical context and current standards." (^[19] www.sei.com). Critical sections were auto-identified and templated, and key compliance checks (e.g. missing criteria) were flagged automatically. While exact time saved is confidential, the firm reported visibly shorter review cycles. *Implications:* This practical pilot confirms that modern AI (when properly integrated) can meaningfully lighten the content creation workload in a Vault-like ecosystem.
- **Merck & Co. (Regulatory Writing AI)** – *Reportedly,* Merck's collaboration with McKinsey led to an AI content generation platform that "reduced first-draft CSR writing time to 80 hours, from 180 hours", with error rates halved (^[58] www.mckinsey.com). *Scenario:* Delivering large clinical reports under tight timelines. *Intervention:* Merck's data scientists trained a domain-specific LLM on internal clinical data. *Outcome:* Complex draft documents (CSRs) were auto-generated to a high-quality standard, slashing manual authoring time almost 2x. *Implications:* This indicates that even critical narrative submissions (e.g. clinical modules) can benefit from GenAI. If similar approaches are applied to regulatory writing in Vault (e.g. auto-drafting Module 1 summaries or RA cover letters), they could substantially reduce cycle time and free experts for higher-value tasks.
- **Ipsen (Submission Content Planning Re-engineering)** – *Scenario:* In a 2014 Veeva R&D summit, Ipsen regulators described their Vault Submissions rollout. They "standardize [d] content management on Vault" and leveraged binders/placeholders to plan their documents ^[43]. *Outcome:* While details are sparse, Ipsen's case implicitly shows that vault-based planning (even pre-AI) can replace complex legacy processes. Ipsen combined Vault's binder templates with company-specific planning, enabling systematic content organization. *Implications:* Even companies not undertaking aggressive process overhaul find that switching to Vault's structured planning yields efficiency. AI can layer on top of that base to suggest improvements (e.g. if Ipsen uploaded multiple dossiers, an AI could refine their binder templates by finding common patterns).
- **Case Example – Translation and Label Submission:** While not publicly documented for Vault specifically, AI-powered translation tools have already been piloted in pharma. For submissions, AI translation can convert modules (or label texts) into other languages with high accuracy (^[54] www.veeva.com). Once translated text is in Vault, regulatory reviewers can focus on minor edits rather than full human translations. A published use-case (not Vault-specific) indicates that "several large biopharmas completed POCs for AI-driven translation" of submission documents, accelerating multilingual filings (^[54] www.veeva.com). This is relevant in RIM because Vault often serves as the repository of multi-language submission assets; AI translation inside Veeva would streamline global submissions planning (eliminating separate translation vendor workflows for high-volume markets). While we lack a specific Vault case for translation, it is a logical extension of submission planning for global filings.

These examples illustrate multiple perspectives: small agile teams (Moderna) vs. large legacy companies (Merck), as well as concrete data (Merck's 56% time cut (^[58] www.mckinsey.com)) and conceptual outcomes (Ipsen, strategic retooling). Together, they support the central thesis: *automation and AI in RIM have genuine, measurable impact on submission speed and quality.*

Discussion: Implications and Future Directions

The integration of AI into Veeva Vault RIM heralds a transformative shift in regulatory operations. Our review of trends, data, and case studies suggests several implications and avenues for the future:

- 1. Accelerated Time to Market:** The primary business driver is speed. By automating rote tasks and leveraging predictive insights, companies can file far more quickly and flexibly. If Vault AI agents eventually enable routine DBL-to-filing in 6–8 weeks or less (^[24] www.mckinsey.com), the payoffs are enormous (e.g. extended patent exclusivity, as McKinsey noted (^[59] www.mckinsey.com)). This could become a key component of commercial strategy: instead of simply planning submissions perfectly, R&D teams will compete on how fast (and how error-free) they can deliver documents. In some areas (like real-time reviews in oncology), regulators themselves are enabling speed, and Vault+AI can help sponsor compliance with those new paradigms.
- 2. Changing Roles in Regulatory Affairs:** As AI takes over repetitive documentation work, human experts will shift toward higher-level oversight, strategy, and creative problem-solving. Regulatory professionals may spend more time on cross-functional decision-making (e.g. designing studies to avoid queries, negotiating with agencies) and less on formatting or quoting text. We might even see new roles emerge, such as "AI Compliance Officer" or "RIM Data Scientist," focused on grooming AI models, ensuring data quality, and bridging the human/AI interface. Training will be critical: users must learn how to "prompt" AI agents effectively (e.g. crafting precise questions to the chatbot), and how to critically evaluate AI outputs.
- 3. Governance and Quality Assurance:** Regulators will demand traceability of AI involvement. This aligns with broader discussions on AI in regulated product development (^[23] www.fda.gov) (^[60] pmc.ncbi.nlm.nih.gov). The FDA's guidance (Jan 2025) on AI in submissions focuses on **model credibility** or "fit for purpose" validation (^[23] www.fda.gov). Applied to Vault RIM, this means sponsors must validate that an AI assistant does not corrupt critical content. Possible safeguards include: (a) limiting AI generation to draft form, never final output without human sign-off; (b) log-based monitoring (every AI action in Vault is logged for audit); © periodic validation exercises (test the AI with known queries to ensure no drift). AI checkpoints may be protocolized similarly to other QA checks (e.g. the "release train engineers" noted in [26] L79-L88] might eventually oversee AI adoption).
- 4. Evolving RIM Standards:** As more data is managed by AI-augmented systems, regulatory standards may adapt. For example, if authorities see high-quality, consistent AI-assisted submissions, they might encourage structured data exchange (e.g. pushing ICH M4 updates for more dynamic eCTD). There is already movement toward "data-rich" submissions (structured content rather than static PDFs) (^[61] www.veeva.com). AI could facilitate that evolution by auto-generating machine-readable content. Furthermore, global harmonization efforts (MHRA, FDA, EMA) may incorporate AI guidance. For instance, the EU AI Act (adopted 2024) classifies FDA/EMA submission tasks as "high-risk" AI applications, implying strict requirements on transparency and oversight (^[62] pmc.ncbi.nlm.nih.gov). Companies will need to track these regulations and ensure Vault AI agents comply (e.g. by disclosing when AI was used in dossier preparation).
- 5. Technological Ecosystem:** The RIM ecosystem will become more interconnected. Already, Vault Submissions can push packages to FDA Gateway and EMA eSub (XEVMPD) through automated publishing (regulatory.veevavault.help). In the future, we may see direct data links between Vault and regulators' systems: for example, Vault's product registration fields could auto-populate portions of the EU's SPOR database (IDMP) or FDA's EP or NAI systems, using AI mapping. On the internal side, Vault RIM may increasingly interface with other Veeva Clouds: Vault Clinical Data (CTMS/CDMS) for real-time updates from trials, Vault Safety for expedited safety reporting, and Vault Quality for deviations that require regulatory notifications. AI will likely be the "glue" that translates between these domains – e.g. parsing a quality incident to determine if it triggers a submission amendment. Interoperability via APIs and AI augmentation is a logical next step.
- 6. Data Requirements and Readiness:** AI is only as good as its data. A recurring theme in our sources is the need for high-quality training data. For regulatory AI, this means maintaining a well-structured Vault history of all past submissions, queries, and outcomes. The Freyr blog emphasizes establishing "a high-quality HA query database" as a prerequisite for advanced automation (^[63] www.veeva.com). Therefore, part of the AI roadmap should be data cleanup and standardization: ensuring all HAQs and HARs are captured in Vault, that documents are consistently tagged, and that sensitive data is anonymized for model training. Data governance will take on added importance.

7. Risk and Mitigation: There are also pessimistic scenarios to consider. Over-reliance on AI without proper validation could introduce errors. Generative models might insert incorrect or outdated regulatory citations (hallucinations), which would then require costly rework. If not carefully managed, AI could also reveal proprietary information (e.g., if an agent was inadvertently exposed to external data in training). Therefore, sponsors must remain vigilant. Human oversight is non-negotiable; systems must be designed so that an AI "suggestion" can never bypass a human check. Furthermore, training models on copyrighted or confidential data (like previously submitted FDA documents) might raise legal issues. These are active research areas in AI governance (^[64] pmc.ncbi.nlm.nih.gov) (^[65] pmc.ncbi.nlm.nih.gov).

8. Competitive Dynamics: As large companies capture the benefits of Vault AI, best practices will be needed industry-wide. Smaller biotechs and generics producers can also benefit if AI agents become a standard Vault offering (not just niche custom developments). Veeva's usage-based AI pricing is intended to democratize adoption (^[66] ir.veeva.com). Over time, consulting firms and Veeva partners will likely offer pre-packaged "Vault AI for RIM" modules (for example, a HAQ-answering assistant or a submission-planning bot). The marketplace may see startups focusing on narrow RIM AI use cases (as we already see with regulatory intelligence services). Ultimately, companies that build internal AI competency in regulatory affairs will have an edge: they can iterate on AI prompts and models to fine-tune for their products and regions.

The future of AI in RIM looks promising. We anticipate a **paradigm shift** where regulatory submissions become more dynamic and data-driven. Imagine submitting portions of an NDA as updates rather than one big packet, with AI continuously checking compliance. Picture an AI avatar running pre-submission mock audits, highlighting weak arguments before a dossier ever goes out. These scenarios are becoming feasible as Vault and AI converge. The industry is rapidly moving in this direction – indeed, participants at the October 2025 Veeva MedTech Summit noted that "AI agents will yield great productivity for life sciences". Our analysis suggests that regulatory teams who embrace these technologies can transition from reactive filing to strategic planning – ultimately bringing therapies to patients faster and more efficiently.

Conclusion

In summary, automating submissions planning and HA correspondence with AI in the context of Veeva Vault RIM has the potential to revolutionize regulatory affairs. By building on Vault's robust data model and workflow engine with advanced AI, companies can substantially cut filing timelines and manpower requirements. Key takeaways from this report include:

- **Proven ROI of Faster Submissions:** Financial analyses show enormous returns (tens of millions in NPV) for each month shaved off submission timelines (^[5] www.mckinsey.com). AI can help capture that ROI by eliminating manual delays and reducing errors that cause rejections or resubmissions.
- **Existing RIM Foundation:** Veeva Vault RIM already provides many automation features (content plans, query tracking) that solve baseline inefficiencies (regulatory.veevavault.help) (regulatory.vaultbasics.help). These should be fully leveraged first (e.g. enforce strict use of content plans, train staff on HA interactions) before layering on AI.
- **Complementary AI Use-Cases:** The highest value AI applications in RIM are those that address known bottlenecks: drafting repetitive content (reports, label text), mining regulatory intelligence, and replying to queries. We identified several concrete scenarios where AI has demonstrated value (^[58] www.mckinsey.com) (^[13] www.freyrdigital.com). Organizations should prioritize these use cases for pilot projects, measuring their impact on cycle time and quality.
- **Case Study Lessons:** Real-world examples (Moderna, Merck, etc.) provide blueprints. A "small team with AI/Automation" can scale like a large one; conversely, a large firm that leverages AI can make common tasks routine. Capturing and reusing organizational knowledge is crucial – Vault's central repository is an excellent platform to do this, especially when augmented by AI search.
- **Challenges to Address:** Key challenges remain around AI trust, data readiness, and change management. The risk of AI hallucinations means all AI outputs must be validated; this should be built into SOPs. Data hygiene (complete, accurate Vault records) is a prerequisite. Both ICH (electronic submission guidelines) and regulatory agencies' AI guidance must be followed to ensure compliance. Continuous oversight (e.g. internal audits of AI usage) will be needed at least during initial adoption.

- **Future Outlook:** The trajectory is clear: every aspect of the submission lifecycle is becoming data-driven and automated. Veeva's planned AI roadmap suggests that "RIM+AI" will become the new standard by 2026–2027. Regulatory authorities, too, are evolving (e.g. FDA's AI frameworks (^[23] www.fda.gov)). In the near future, we may see fully paperless filings, real-time regulator–sponsor data exchanges, and predictive regulatory intelligence as routine. Achieving that will require a human-in-the-loop approach; expert judgment and ethical oversight remain indispensable.

For life sciences companies, failure to leverage these technologies will look like falling behind in an increasingly digital regulatory environment. The evidence suggests that careful, strategic integration of AI into Veeva Vault RIM can yield huge dividends in speed, consistency, and insight. This report provides a deep dive into both the promise and the caution points, giving regulatory leaders a roadmap for bringing AI into their submission planning and HA correspondence workflows. By iterating on the use cases and best practices outlined here, companies can shape a future in which regulatory teams spend less time on paperwork and more time on science and strategy – ultimately accelerating patient access to life-saving therapies.

External Sources

- [1] <https://www.veeva.com/customer-stories/moderna-improves-health-authority-query-management-with-veeva-vault-rim/#:~:Like%...>
- [2] <https://www.freysolutions.com/blog/revolutionize-regulatory-publishing-submission-with-automation#:~:For%2...>
- [3] <https://www.veeva.com/eu/blog/picking-up-the-pace-how-rim-agility-enables-proactive-compliance/#:~:Next%...>
- [4] <https://www.veeva.com/products/vault-submissions/#:~:15...>
- [5] <https://www.mckinsey.com/industries/life-sciences/our-insights/rewiring-pharmas-regulatory-submissions-with-ai-and-zero-based-design#:~:subst...>
- [6] <https://www.mckinsey.com/industries/life-sciences/our-insights/rewiring-pharmas-regulatory-submissions-with-ai-and-zero-based-design#:~:In%20...>
- [7] <https://www.mckinsey.com/industries/life-sciences/our-insights/rewiring-pharmas-regulatory-submissions-with-ai-and-zero-based-design#:~:Gen%2...>
- [8] <https://www.mckinsey.com/industries/life-sciences/our-insights/rewiring-pharmas-regulatory-submissions-with-ai-and-zero-based-design#:~:Block...>
- [9] <https://www.veeva.com/resources/vault-rim-submissions-content-plan-hierarchy-viewer/#:~:Submit...>
- [10] <https://www.veeva.com/customer-stories/moderna-improves-health-authority-query-management-with-veeva-vault-rim/#:~:%2A%2...>
- [11] <https://www.veeva.com/customer-stories/moderna-improves-health-authority-query-management-with-veeva-vault-rim/#:~:Now%2...>
- [12] <https://www.veeva.com/blog/emerging-ai-proof-of-concepts-for-accelerating-submission-timelines/#:~:It%E2...>
- [13] <https://www.freydigital.com/blog/the-ai-advantage-in-health-authority-queries-management#:~:NLP,T...>
- [14] <https://www.freydigital.com/blog/the-ai-advantage-in-health-authority-queries-management#:~:,vali...>
- [15] <https://www.freydigital.com/blog/the-ai-advantage-in-health-authority-queries-management#:~:AI%20...>
- [16] <https://www.mckinsey.com/industries/life-sciences/our-insights/rewiring-pharmas-regulatory-submissions-with-ai-and-zero-based-design#:~:submi...>
- [17] <https://www.freydigital.com/blog/the-ai-advantage-in-health-authority-queries-management#:~:1,ins...>

- [18] <https://www.mckinsey.com/industries/life-sciences/our-insights/rewiring-pharmas-regulatory-submissions-with-ai-and-zero-based-design#:~:Autom...>
- [19] <https://www.sei.com/case-study/streamline-pharmaceutical-regulatory-submission-with-generative-ai#:~:subm...>
- [20] <https://ir.veeva.com/investors/news-and-events/latest-news/press-release-details/2025/Announcing-Veeva-AI/default.aspx#:~:PLEAS...>
- [21] <https://ir.veeva.com/news/news-details/2025/Veeva-AI-Agents-to-Be-Released-Across-All-Veeva-Applications/default.aspx#:~:Plan...>
- [22] <https://www.veeva.com/products/vault-submissions/#:~:...>
- [23] <https://www.fda.gov/news-events/press-announcements/fda-proposes-framework-advance-credibility-ai-models-used-drug-and-biological-product-submissions#:~:A%20k...>
- [24] <https://www.mckinsey.com/industries/life-sciences/our-insights/rewiring-pharmas-regulatory-submissions-with-ai-and-zero-based-design#:~:Horiz...>
- [25] <https://www.mckinsey.com/industries/life-sciences/our-insights/rewiring-pharmas-regulatory-submissions-with-ai-and-zero-based-design#:~:Almos...>
- [26] <https://www.veeva.com/eu/blog/picking-up-the-pace-how-rim-agility-enables-proactive-compliance/#:~:The%2...>
- [27] <https://www.veeva.com/eu/blog/picking-up-the-pace-how-rim-agility-enables-proactive-compliance/#:~:Aroun...>
- [28] <https://www.veeva.com/eu/blog/picking-up-the-pace-how-rim-agility-enables-proactive-compliance/#:~:In%20...>
- [29] <https://www.veeva.com/eu/blog/picking-up-the-pace-how-rim-agility-enables-proactive-compliance/#:~:Separ...>
- [30] <https://www.veeva.com/eu/blog/picking-up-the-pace-how-rim-agility-enables-proactive-compliance/#:~:Overa...>
- [31] <https://www.veeva.com/products/vault-submissions/#:~:Veeva...>
- [32] <https://www.veeva.com/products/vault-submissions/#:~:Strea...>
- [33] <https://www.veeva.com/products/vault-submissions/#:~:Submi...>
- [34] <https://www.veeva.com/customer-stories/moderna-improves-health-authority-query-management-with-veeva-vault-rim/#:~:Despi...>
- [35] <https://www.veeva.com/customer-stories/moderna-improves-health-authority-query-management-with-veeva-vault-rim/#:~:With%...>
- [36] <https://www.freydigital.com/blog/the-ai-advantage-in-health-authority-queries-management#:~:Autom...>
- [37] <https://app.submitly.ai/contracts/category/48000000#:~:Care%...>
- [38] <https://www.veeva.com/products/vault-submissions/#:~:Read%...>
- [39] <https://www.mckinsey.com/industries/life-sciences/our-insights/rewiring-pharmas-regulatory-submissions-with-ai-and-zero-based-design#:~:Block...>
- [40] <https://iernal.ai/calculators/pharmaceutical-regulatory-submission-calculator#:~:AI%20...>
- [41] <https://www.veeva.com/blog/emerging-ai-proof-of-concepts-for-accelerating-submission-timelines/#:~:Autom...>
- [42] <https://ir.veeva.com/investors/news-and-events/latest-news/press-release-details/2025/Announcing-Veeva-AI/default.aspx#:~:use...>
- [43] <https://www.mckinsey.com/industries/life-sciences/our-insights/rewiring-pharmas-regulatory-submissions-with-ai-and-zero-based-design#:~:Categ...>
- [44] <https://www.mckinsey.com/industries/life-sciences/our-insights/rewiring-pharmas-regulatory-submissions-with-ai-and-zero-based-design#:~:Horiz...>
- [45] <https://www.freydigital.com/blog/the-ai-advantage-in-health-authority-queries-management#:~:%2A%2...>

- [46] <https://www.freydigital.com/blog/the-ai-advantage-in-health-authority-queries-management#:~:AI%20...>
- [47] <https://ir.veeva.com/investors/news-and-events/latest-news/press-release-details/2025/Announcing-Veeva-AI/default.aspx#:~:worki...>
- [48] <https://ir.veeva.com/investors/news-and-events/latest-news/press-release-details/2025/Announcing-Veeva-AI/default.aspx#:~:the%2...>
- [49] <https://ir.veeva.com/news/news-details/2025/Veeva-AI-Agents-to-Be-Released-Across-All-Veeva-Applications/default.aspx#:~:hoste...>
- [50] <https://pmc.ncbi.nlm.nih.gov/articles/PMC12675821#:~:Recen...>
- [51] <https://www.veeva.com/products/vault-submissions/#:~:Watch...>
- [52] <https://www.freysolutions.com/blog/revolutionize-regulatory-publishing-submission-with-automation#:~:%2A%2...>
- [53] <https://www.veeva.com/blog/emerging-ai-proof-of-concepts-for-accelerating-submission-timelines#:~:,Prat...>
- [54] <https://www.veeva.com/blog/emerging-ai-proof-of-concepts-for-accelerating-submission-timelines#:~:,Strea...>
- [55] <https://www.veeva.com/blog/emerging-ai-proof-of-concepts-for-accelerating-submission-timelines#:~:,Lever...>
- [56] <https://www.mckinsey.com/industries/life-sciences/our-insights/rewiring-pharmas-regulatory-submissions-with-ai-and-zero-based-design#:~:McKin...>
- [57] <https://www.veeva.com/customer-stories/moderna-improves-health-authority-query-management-with-veeva-vault-rim#:~:Since...>
- [58] <https://www.mckinsey.com/industries/life-sciences/our-insights/rewiring-pharmas-regulatory-submissions-with-ai-and-zero-based-design#:~:Gen%2...>
- [59] <https://www.mckinsey.com/industries/life-sciences/our-insights/rewiring-pharmas-regulatory-submissions-with-ai-and-zero-based-design#:~:savin...>
- [60] <https://pmc.ncbi.nlm.nih.gov/articles/PMC12675821#:~:In%20...>
- [61] <https://www.veeva.com/eu/blog/picking-up-the-pace-how-rim-agility-enables-proactive-compliance#:~:direc...>
- [62] <https://pmc.ncbi.nlm.nih.gov/articles/PMC12675821#:~:landm...>
- [63] <https://www.veeva.com/blog/emerging-ai-proof-of-concepts-for-accelerating-submission-timelines#:~:,have%...>
- [64] <https://pmc.ncbi.nlm.nih.gov/articles/PMC12675821#:~:the%2...>
- [65] <https://pmc.ncbi.nlm.nih.gov/articles/PMC12675821#:~:solut...>
- [66] <https://ir.veeva.com/investors/news-and-events/latest-news/press-release-details/2025/Announcing-Veeva-AI/default.aspx#:~:The%2...>

IntuitionLabs - Industry Leadership & Services

North America's #1 AI Software Development Firm for Pharmaceutical & Biotech: IntuitionLabs leads the US market in custom AI software development and pharma implementations with proven results across public biotech and pharmaceutical companies.

Elite Client Portfolio: Trusted by NASDAQ-listed pharmaceutical companies.

Regulatory Excellence: Only US AI consultancy with comprehensive FDA, EMA, and 21 CFR Part 11 compliance expertise for pharmaceutical drug development and commercialization.

Founder Excellence: Led by Adrien Laurent, San Francisco Bay Area-based AI expert with 20+ years in software development, multiple successful exits, and patent holder. Recognized as one of the top AI experts in the USA.

Custom AI Software Development: Build tailored pharmaceutical AI applications, custom CRMs, chatbots, and ERP systems with advanced analytics and regulatory compliance capabilities.

Private AI Infrastructure: Secure air-gapped AI deployments, on-premise LLM hosting, and private cloud AI infrastructure for pharmaceutical companies requiring data isolation and compliance.

Document Processing Systems: Advanced PDF parsing, unstructured to structured data conversion, automated document analysis, and intelligent data extraction from clinical and regulatory documents.

Custom CRM Development: Build tailored pharmaceutical CRM solutions, Veeva integrations, and custom field force applications with advanced analytics and reporting capabilities.

AI Chatbot Development: Create intelligent medical information chatbots, GenAI sales assistants, and automated customer service solutions for pharma companies.

Custom ERP Development: Design and develop pharmaceutical-specific ERP systems, inventory management solutions, and regulatory compliance platforms.

Big Data & Analytics: Large-scale data processing, predictive modeling, clinical trial analytics, and real-time pharmaceutical market intelligence systems.

Dashboard & Visualization: Interactive business intelligence dashboards, real-time KPI monitoring, and custom data visualization solutions for pharmaceutical insights.

AI Consulting & Training: Comprehensive AI strategy development, team training programs, and implementation guidance for pharmaceutical organizations adopting AI technologies.

Contact founder Adrien Laurent and team at <https://intuitionlabs.ai/contact> for a consultation.

DISCLAIMER

The information contained in this document is provided for educational and informational purposes only. We make no representations or warranties of any kind, express or implied, about the completeness, accuracy, reliability, suitability, or availability of the information contained herein.

Any reliance you place on such information is strictly at your own risk. In no event will IntuitionLabs.ai or its representatives be liable for any loss or damage including without limitation, indirect or consequential loss or damage, or any loss or damage whatsoever arising from the use of information presented in this document.

This document may contain content generated with the assistance of artificial intelligence technologies. AI-generated content may contain errors, omissions, or inaccuracies. Readers are advised to independently verify any critical information before acting upon it.

All product names, logos, brands, trademarks, and registered trademarks mentioned in this document are the property of their respective owners. All company, product, and service names used in this document are for identification purposes only. Use of these names, logos, trademarks, and brands does not imply endorsement by the respective trademark holders.

IntuitionLabs.ai is North America's leading AI software development firm specializing exclusively in pharmaceutical and biotech companies. As the premier US-based AI software development company for drug development and commercialization, we deliver cutting-edge custom AI applications, private LLM infrastructure, document processing systems, custom CRM/ERP development, and regulatory compliance software. Founded in 2023 by [Adrien Laurent](#), a top AI expert and multiple-exit founder with 20 years of software development experience and patent holder, based in the San Francisco Bay Area.

This document does not constitute professional or legal advice. For specific guidance related to your business needs, please consult with appropriate qualified professionals.

© 2025 IntuitionLabs.ai. All rights reserved.